

Social Media as Windows on the Social Life of the Mind

Cosma Rohilla Shalizi

Statistics Department, Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213 USA

Abstract

This is a programmatic paper, marking out two directions in which the study of social media can contribute to broader problems of social science: understanding cultural evolution and understanding collective cognition. Under the first heading, I discuss some difficulties with the usual, adaptationist explanations of cultural phenomena, alternative explanations involving network diffusion effects, and some ways these could be tested using social-media data. Under the second I describe some of the ways in which social media could be used to study how the social organization of an epistemic community supports its collective cognitive performance.

Let me begin by considering two¹ senses in which we might speak of human thought as being “social”, and how they might orient the study of social information processing and social media.

The first sense is a common-place of many schools in the social sciences and humanities: our thought relies on the cultural transmission of cognitive tools. Every individual thinker, no matter how innovative or even lonely they may be, depends crucially on a vast array of cognitive tools (concepts, procedures, languages, assumptions, values, ...) which they did not devise themselves, and could not have devised for themselves. Instead they inherited these cognitive tools from interacting with other people, who for the most part themselves did not invent them. (Dewey 1927; Vygotsky 1934/1986; Popper 1945; Balkin 1998)² (Whether this dependence on tradition is a logical necessity, or merely a reflection

of our peculiar bounded rationality and bounded lifespan, is a deep question, fortunately not relevant here.) While individual thinkers invent and discover, it is nonetheless true that innovations are typically refined, extended and perfected by groups, and that it is very rare indeed for highly developed concepts and ideas to emerge from a single, isolated thinker, rather than from a process of interaction (Toulmin 1972; Kitcher 1993; Collins 1998; Ziman 2000).

The branches of social science for which these facts are common-places have largely developed them philosophically (Toulmin 1972; Turner 2002), or qualitatively (Vygotsky 1934/1986; 1978; Balkin 1998; Mercer 2000) or even ethnographically (Luria 1976; Hutchins 1995). (But see (Lupia, McCubbins, & Popkin 2000).) In part this has been for reasons of cultural and intellectual politics, as the relevant scholars have tended to fall on the “interpretation” rather than “explanation” side of the divide in the social sciences (Sperber 1996), so that attention to the social nature of thought often goes along with more or less pronounced hostility to quantitative and computational modeling (e.g. Hutchins; Mercer). This supposed opposition is thoroughly mis-guided (Frawley 1997), but it is not likely that anyone will be argued out of it any time soon.

More promisingly, however, one *good* reason for developing this idea through small-scale qualitative studies has been that it was impossible to gather *relevant* data, suitable for quantitative analysis, on any large scale. With the rise of social media, however, many people are, for their own purposes, generating exactly this kind of data for us — traces of their communicative interactions as they work out their thoughts about matters of common concern. They are doing so on a wide range of subjects, under a wide range of different institutional mechanisms which structure their interactions in many different ways, creating natural sources of variation which the social scientist can try to exploit to learn more about the effects of subject matter, of communicative structure, and of other factors on cultural dynamics, and perhaps ultimately even on innovation and discovery. The next section points out some of the outstanding problems and methodological pitfalls of this area.

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Of course, people think a lot about their own and others’ social interactions, and a big use of social media is sharing these thoughts. But in this social media are no different from any other form of human, or for that matter primate, association.

² “[K]nowledge is a function of association and communication; it depends upon tradition, upon tools and methods socially transmitted, developed and sanctioned. Faculties of effectual observation, reflection and desire are habits acquired under the influence of the culture and institutions of society, not ready-made inherent powers” (Dewey 1927, p. 158). Cf. (Popper 1945, ch. 23–24).

The *other* important sense in which human thought can be “social” is that it seems to make sense to regard at least some human social institutions as, themselves, information-processing systems, engaging in computations which cannot be localized to representations in the mind of any one of their members. On large scales, market economies, corporations and other bureaucracies, scientific disciplines, and democratic polities all have something of this collective information-processing character. Knowing how they accomplish this would be deeply rewarding, and, if that understanding can be used to make them work better, of profound economic and political importance. A frontal assault on this problem, as represented by one of those grand institutions, is unlikely to succeed (though it may be a magnificent failure). Fortunately, social information processing also occurs in much humbler institutions, such as tagging systems and collaborative filtering, where issues of data collection and even experimental manipulation are much more manageable, and where we might hope to learn more, before tackling the fundamental problems of social science. I will lay out some of what should be on the agenda of the study of social information processing, in particular points of contact with machine learning.

Cultural Evolution

“Culture is the precipitate of cognition and communication in a human population” (Sperber 1996). That is, cultural traits — beliefs, practices, habits, conventions, expressions, norms — are not just ones which are common across a population, but ones which are spread across a population because its members communicate with one another. (Knowing that it’s painful to look at the sun directly is not cultural; knowing that the direction in which the sun rises is called “east” is cultural.) Cultural phenomena are thus emergent, the result of the communicative interaction of cognitive agents. If we are to understand how cultures work, we need to understand something about both parts, the internal cognitive mechanisms and the effects of different patterns of interaction. Social media offer a window into the communicative part of the problem of unrivaled clarity and breadth. This is extremely exciting, but in looking through this window we should bear in mind some methodological difficulties to interpreting the view through this window.

It is a common-place observation that there are strong relationships between cultural traits and social attributes; that different social groups accept and transmit different bits of culture. Most attempts to explain this from within the social sciences (emphatically including historical materialism (Elster 1985; Cohen 2000) and its variants) argue that this is due to some causal influence of social organization on the *content* of culture. (“Social being determines consciousness” (Marx & Engels 1847/1947) — or, once the Hegelian gas has been released, social life shapes thought.) In these views, culture varies with social po-

sition *because* the former is adapted to the latter, or reflects it, or expresses it.

It is natural for us, as beings acutely sensitive to nuances of cultural meanings, to try to explain cultural differences by trying to explain the *content* of widely-shared, cultural representations. It is natural to suppose that, say, one news story rises to the top of a social aggregation system because it is *more interesting* than other stories which did not. Such explanations are even valid a lot of the time. It is nonetheless important, as a point of methodological hygiene, to develop ways of telling when some bit of culture succeeds in propagating because its content fits its circumstances, if only because, being creatures acutely sensitive to nuances of cultural meanings, it is far too easy for us to spin such stories no matter what the truth might be. Lieberman (2000) points out that many widely-accepted explanations of trends in fashions, children’s names, etc., cannot possibly be right (because, e.g., the trend pre-dates or is more widely spread than the supposed cause), and that these are instead better explained by purely internal mechanisms of the respective fields. In biology, adaptive and non-adaptive evolution are demarcated by means of *neutral models*. These are models of the genetic changes which would be expected due to reproductive mechanisms and chance alone, all genetic variants being assumed to be “adaptively neutral”, i.e., of equal fitness. Only when actual populations depart markedly from the predictions of neutral models can adaptation be (reliably) inferred (Nitecki & Hoffman 1987; Harvey & Pagel 1991). Before the student of social media, or other cultural media, can start explaining phenomena by reference to content, they need to check that there actually is something to be explained.

A highly simplistic model may make this point more concrete. Consider a network in which people have two binary traits, one of which is stable (we may think of this as “class” or “race” or some similar status), and the other is changeable (think of fashions, or political opinions). Assume that the network is assortative on the stable, social-type trait, so that people are more likely to be linked to others of the same type than those of a different type. Such “assortativity” or “homophily” is observed in many, perhaps most social networks, often on such stable social-status type variables (McPherson, Smith-Lovin, & Cook 2001; Newman 2003). Now assign the cultural trait to people uniformly and independently of their social trait (or anything else). Initially, then, there will be no correlation between social and cultural traits, and no assortativity based on culture.

We might expect such correlations to appear if the process of cultural transmission and retention is biased — if, say, certain cultural values only make sense for those in certain social positions. In that case, we would expect to find a growing “fit” between cultural and social variables, as the former adapt to the latter. But by this point you will not be surprised to learn that *neutral* transmission processes can also induce such cor-




Figure 1: Neutral copying induces correlations between social and cultural traits in assortative networks. The graph has 100 nodes, randomly divided between two social types (equally probable), and a binary-valued cultural trait (initially equally probable). Edges between nodes of the same type occur with probability p_1 , those between different types have probability p_2 . At each time step, a random node copies the cultural trait of a random neighbor. Horizontal axis: time. Vertical axis: χ^2 statistic for the correlation between the social and cultural variables. Black line: behavior of an assortative network ($p_1 = 0.09$, $p_2 = 0.01$, assortativity coefficient (Newman 2003) of realized graph $r = 0.80$). Note the eventual decline of χ^2 as the network moves towards a homogeneous equilibrium; in the very long run it will reach 0. Grey line: behavior of a non-assortative network ($p_1 = p_2 = 0.05$, $r = 0.045$).

relations. To be specific, let’s implement the “voter model” (Liggett 1985): at each discrete time step, a node is chosen uniformly at random, independently of past and future choices. This node chooses a neighbor (again uniformly and independently), and copies its value of the cultural trait.

Clearly, in a connected network, there are two absorbing states, which are culturally homogeneous, and eventually the network must settle into one or the other of them, but the time it takes to do so will typically be quite long (Sood & Redner 2005). In the meanwhile, if the network is socially assortative, numerical experiments (Fig. 1) show that the social and the cultural traits tend to become correlated during a long “meta-stable” period.³ If I’d said that the social types

³One could say that the cultural trait of a node is still “in the final analysis” determined by its social type, but only with the proviso that the over-all structure of the network “screens off” the latter, rendering it causally irrelevant

were “lower class” and “middle class”, and the cultural traits “likes black velvet paintings” and “likes black and white photographs”, the temptation to explain the correlation by content would be overwhelming (Bourdieu 1984). Nonetheless, which way the correlation went would be a matter of pure chance, or more exactly of the reinforcement and amplification of small fluctuations, though *some* such pattern forms with high probability. (This contrast between long- and medium- run behavior is not uncommon in self-reinforcing network processes (Pemantle & Skyrms 2004).) The strength of the dynamically-induced correlations depends on the assortativity of the social network; if it is not assortative, then the correlations between social and cultural traits only rarely rise above the levels to be expected by chance (Fig. 1).

Social scientists interested in communications have appreciated for a long time that network structure is very important to *how* information flows through a social group (Katz & Lazarsfeld 1955; Huckfeldt, Johnson, & Sprague 2004), but they have not, so far as I know, realized that it can create just the kind of correlation that seems to cry out for an explanation by content. In fact, the real situation is somewhat worse than this, because it really isn’t a given that people change because of interacting with their neighbors. It could well be that people have neighbors who are similar to themselves, and so they all respond similarly to common exogenous causes, without any direct interactions (Steglich, Snijders, & Pearson 2004).⁴ If one thinks of trying to explain why certain users prefer certain kinds of news stories, for example, one must account not only assortativity, but also for common exposure to some outside news source. — None of this, incidentally, requires that people *actually* make decisions randomly, but only that the reasons which lead them to their decisions are effectively unpredictable from the other variables in the system.

The moral is not that these kinds effect explain *all* correlations between social and cultural traits, or even between different cultural traits. Rather, it shows that a neutral explanation is logically possible. To support an adaptive explanation of a correlation, then, one must show some way in which the neutral model is not adequate to the data. For example, additional experiments (not shown) indicate that, if I take the model simulated in Fig. 1 and break the graph into communities (following Newman & Girvan (2003)), then social type and cultural traits are *conditionally* independent, given community membership, even in strongly-assortative networks. This conditional independence does *not* hold when different social types have differing biases for or against various cultural traits. Only when we have found and verified such discrepancies between

(Galles & Pearl 1997).

⁴This possibility seems to confound the claims of the recent, and widely-publicized, study of the spread of obesity in a social network (Christakis & Fowler 2007).

our data and the predictions of a good neutral model can we say that the adaptive explanation has passed a severe test and truly has evidence in its support (Mayo 1996).

Collective Cognition

It's been recognized since the 1930s that market economies are "collective calculating devices" (Lange & Taylor 1938; Hayek 1948). A market-clearing allocation of good and services is simply too big for anyone to grasp, let alone find. Instead it is the process of exchange itself which adaptively finds and implements this allocation.⁵ This is an example of what we might call *collective cognition*, by analogy to the classical (Mancur Olson 1971) "collective action". Similarly, the problems of designing policies for governments are largely beyond the scope of what anyone can actually do, but not beyond the scope of democratic deliberation, which reduces the problem from solving for the optimal policy in one stroke, to criticizing and improving policies piecemeal (Braybrooke & Lindblom 1963), in light of the information and ideas of many participants. (Popper 1945; Lindblom 1965; Ober 2005) (Historically, democratic decision-making has been associated with more social power than other forms of government (McNeill 1982), but the causality is unclear.) Similar remarks apply to bureaucratic organizations, such as corporations, and to scientific disciplines.

It is notable that modern societies are vastly better at collective cognition than earlier ones. The degree of organization, and its precision, which we take for granted would have been astonishing for even the inhabitants of the most advanced societies c. 1600, to say nothing of c. 100. Historians have explored some of the technical and institutional underpinnings of these organizational revolutions (McNeill 1982; Beniger 1986; Yates 1989), but at a deeper level we have little idea *why* this is so, or why what we do works (when it does work). This makes it harder to *improve* the functioning of our institutions for collective cognition. Economic theories of mechanism design attempt to do so, but largely address the problem of *motivating* people to act in certain ways, rather than of how to figure out what the right action is (Miller 1992).

These are all very large themes indeed, of course, and it might seem grandiose to even mention them in this context. I am not suggesting that studying social media will give us the key to all organization technologies. What it can do, however, is give us a set of case studies where, on a much humbler level, people are nonetheless engaged in social information processing and collective cognition. Just as no one market participant decides on or represents the over-all market allocation, and no one scholar ever grasps more than a small portion of what is known about conic sections or cellular slime molds,

the movies or bookmarks which get recommended by collaborative filtering services are the emergent products of the interactions of many participants (Lerman 2007). What social media offer us, again, is the possibility to automatically collect large-scale data on such phenomena, combined with a clear understanding of the interaction structure (or at least a lot of it), as well as much of the external circumstances and the goals of the group. We can thus begin, at least at a small scale, to begin building and systematically testing theories which explain how social information processing and collective cognition succeed when they do.

It might be thought that the theoretical explanation is rather simple, and goes (currently) under the name of "the wisdom of crowds" (Surowiecki 2004): individuals make noisy guesses, which on average are unbiased and uncorrelated, so simple averaging leads to convergence on the appropriate answer. Taken seriously, this explanation implies that our economy, our sciences and our politics manage to work *despite* their social organization, that science (for example) would progress much faster if scientists did not collaborate, did not read each others' papers, etc. While every scientist feels this way occasionally, it is hard to take seriously. Clearly, there has to be an explanation for the success of social information processing *other than* averaging uncorrelated guesses, something which can handle, and perhaps even exploit, statistical dependence between decision makers.

A particularly interesting line of attack on these problems is suggested by the analogy with ensemble methods in machine learning. As Domingos (1999) has pointed out, the success of these methods seems to confound naive interpretations of Occam's Razor, in much the same way that the success of social information processing confounds the simple "wisdom of the crowds" story. Ensemble methods, in which large numbers of low-capacity classifiers or predictors (e.g., shallow classification trees) are combined, effectively create a single model of what appears to be very high capacity, and so they appear to be nothing but an invitation to over-fitting. Worse, typically ensemble methods such as boosting (Hastie, Tibshirani, & Friedman 2001), bagging (Breiman 1996) and mixtures of experts (Jacobs 1997) create *correlated* low-level predictors, so that the simple average-the-crowd story is inapplicable. In fact, it is precisely because the component predictors are correlated, but not identical, that the *actual* capacity of the ensemble is much smaller than its *apparent* capacity.

A similar result holds for cooperative problem-solving (Hong & Page 2004). Under mild conditions, it can be shown that a large group of "weak" heuristic problem-solvers, whose performance in isolation is only slightly better than random search, will actually out-perform a similarly-sized group of "strong" heuristics, ones whose average performance in isolation is much better. One of those conditions, however, is that the problem-solvers must be able to communicate with each other, making their candidate solutions strongly dependent rather

⁵On the formal computational power of market-like systems, see (Walsh *et al.* 2003).

than uncorrelated. There is good evidence that this beneficial effect of heuristic diversity and communication is actually seen in the cognitive performance of human groups (Page 2007). This suggests a very promising direction for research on social information processing, namely to use the mathematical techniques of statistical learning theory to establish bounds on the performance of suitable sorts of ensemble-learners and group problem-solvers, and see how close actual social information processing systems come to attain those bounds, and how the latter could be improved by changes to their architectures.

Both ensemble methods and the Hong & Page results on diverse heuristics posit relatively simple forms of “social” organization, such as direct averaging, or passing a problem to the next person able to improve on the current solution. There is every reason to think, however, that the optimal form of organization will actually depend on the structure of the problem being solved. (Cf. Braybrooke & Lindblom (1963) on how the social organization of policy analysts serves their cognitive strategy of “disjointed incrementalism.”) In particular, coordination over time is not an issue in ensemble methods, and handled by assumption in the Hong & Page model, but extremely important in real-world systems for social information processing and collective cognition.

This suggests a final line of research, one which draws together ideas from distributed systems, economics and statistical mechanics. Experience with distributed systems shows that often the hardest part of their design is ensuring coordination over time, and that failure to do so can lead to all manner of unwanted behavior, in particular to wild oscillations and/or locking into deeply undesirable configurations (Lynch 1996). In fact, the failure modes of distributed systems are strongly reminiscent of the pathologies of economic (Chamley 2004) and statistical-mechanical (Young 1998) models of social learning, when they are placed in suitable (that is, unsuitable) situations. Designing, or reforming, a system for computer-mediate social information processing is at once a problem of distributed algorithm design and a problem of mechanism design, and they two modes or aspects should inform one another, as well as empirical results about what actually happens when real human beings use different systems for different tasks.

Acknowledgments Thanks to P. Agre, P. Domingos, C. Genovese, K. Lee, S. Page, W. Tozier, E. Smith, N. Snod, and the participants of the 2002 workshop on collective cognition and distributed intelligence at the Santa Fe Institute for valuable discussions, and to K. Klinkner for many reasons (including valuable discussions).

References

- Balkin, J. M. 1998. *Cultural Software: A Theory of Ideology*. New Haven, Connecticut: Yale University Press.
- Beniger, J. 1986. *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge, Massachusetts: Harvard University Press.
- Bourdieu, P. 1984. *Distinction: A Social Critique of the Judgement of Taste*. Cambridge, Massachusetts: Harvard University Press.
- Braybrooke, D., and Lindblom, C. E. 1963. *A Strategy of Decision: Policy Evaluation as a Social Process*. Glencoe, Illinois: The Free Press.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* 24:123–140.
- Chamley, C. 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge, England: Cambridge University Press.
- Christakis, N. A., and Fowler, J. H. 2007. The spread of obesity in a large social network over 32 years. *The New England Journal of Medicine* 357:370–379.
- Cohen, G. A. 2000. *Karl Marx’s Theory of History: A Defense*. Princeton, New Jersey: Princeton University Press, second edition.
- Collins, R. 1998. *The Sociology of Philosophies: A Global Theory of Intellectual Change*. Cambridge, Massachusetts: Harvard University Press.
- Dewey, J. 1927. *The Public and Its Problems*. New York: Henry Holt.
- Domingos, P. 1999. The role of Occam’s Razor in knowledge discovery. *Data Mining and Knowledge Discovery* 3:409–425.
- Elster, J. 1985. *Making Sense of Marx*. Cambridge, England: Cambridge University Press.
- Frawley, W. D. 1997. *Vygotsky and Cognitive Science: Language and the Unification of the Social and Computational Mind*. Cambridge, Massachusetts: Harvard University Press.
- Galles, D., and Pearl, J. 1997. Axioms of causal relevance. *Artificial Intelligence* 97:9–43.
- Harvey, P. H., and Pagel, M. D. 1991. *The Comparative Method in Evolutionary Biology*. Oxford: Oxford University Press.
- Hastie, T.; Tibshirani, R.; and Friedman, J. 2001. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer-Verlag.
- Hayek, F. A. 1948. *Individualism and Economic Order*. Chicago: University of Chicago Press.
- Hong, L., and Page, S. E. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* 101:16385–16389.
- Huckfeldt, R.; Johnson, P. E.; and Sprague, J. 2004. *Political Disagreement: The Survival of Diverse Opinions within Communication Networks*. Cambridge, England: Cambridge University Press.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, Massachusetts: MIT Press.

- Jacobs, R. A. 1997. Bias/variance analyses of mixtures-of-experts architectures. *Neural Computation* 9:369–383.
- Katz, E., and Lazarsfeld, P. F. 1955. *Personal influence: the part played by people in the flow of mass communications*. Glencoe, Illinois: Free Press.
- Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.
- Lange, O., and Taylor, F. M. 1938. *On the Economic Theory of Socialism*. Minneapolis: University of Minnesota Press.
- Lerman, K. 2007. Social information processing in social news aggregation. *IEEE Internet Computing* submitted.
- Lieberson, S. 2000. *A Matter of Taste: How Names, Fashions, and Culture Change*. New Haven, Connecticut: Yale University Press.
- Liggett, T. M. 1985. *Interacting Particle Systems*. Berlin: Springer-Verlag.
- Lindblom, C. E. 1965. *The Intelligence of Democracy: Decision Making through Mutual Adjustment*. New York: Free Press.
- Lupia, A.; McCubbins, M. D.; and Popkin, S. L., eds. 2000. *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*. Cambridge, England: Cambridge University Press.
- Luria, A. R. 1976. *Cognitive Development: Its Cultural and Social Foundations*. Cambridge, Massachusetts: Harvard University Press.
- Lynch, N. A. 1996. *Distributed Algorithms*. San Francisco: Morgan Kaufmann.
- Mancur Olson, J. 1971. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, Massachusetts: Harvard University Press, revised edition. First edition, 1965.
- Marx, K., and Engels, F. 1847/1947. *The German Ideology*. New York: International Publishers.
- Mayo, D. G. 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
- McNeill, W. H. 1982. *The Pursuit of Power: Technology, Armed Force and Society since A.D. 1000*. Chicago: University of Chicago Press.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27:415–444.
- Mercer, N. 2000. *Words and Minds: How We Use Language to Think Together*. London: Routledge.
- Miller, G. J. 1992. *Managerial Dilemmas: The Political Economy of Hierarchy*. Cambridge, England: Cambridge University Press.
- Newman, M. E. J., and Girvan, M. 2003. Finding and evaluating community structure in networks. *Physical Review E* 69:026113.
- Newman, M. E. J. 2003. Mixing patterns in networks. *Physical Review E* 67:026126.
- Nitecki, M. H., and Hoffman, A., eds. 1987. *Neutral models in biology*. New York: Oxford University Press.
- Ober, J. 2005. *Athenian Legacies: Essays on the Politics of Going On Together*. Princeton, New Jersey: Princeton University Press.
- Page, S. E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, New Jersey: Princeton University Press.
- Pemantle, R., and Skyrms, B. 2004. Network formation by reinforcement learning: the long and medium run. *Mathematical Social Sciences* 48:315–327.
- Popper, K. R. 1945. *The Open Society and Its Enemies*. London: Routledge.
- Sood, V., and Redner, S. 2005. Voter model on heterogeneous graphs. *Physical Review Letters* 94:178701.
- Sperber, D. 1996. *Explaining Culture: A Naturalistic Approach*. Oxford: Basil Blackwell.
- Steglich, C.; Snijders, T. A. B.; and Pearson, M. 2004. Dynamic networks and behavior: Separating selection from influence. Technical Report 95-2001, Interuniversity Center for Social Science Theory and Methodology, University of Groningen.
- Surowiecki, J. 2004. *The Wisdom of Crowds: Why the Many Are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. New York: Doubleday.
- Toulmin, S. 1972. *Human Understanding: The Collective Use and Evolution of Concepts*. Princeton, New Jersey: Princeton University Press.
- Turner, S. P. 2002. *Brains/Practices/Relativism: Social Theory After Cognitive Science*. Chicago: University of Chicago Press.
- Vygotsky, L. S. 1934/1986. *Thought and Language*. Cambridge, Massachusetts: MIT Press.
- Vygotsky, L. S. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, Massachusetts: Harvard University Press.
- Walsh, W. E.; Yokoo, M.; Hirayama, K.; and Wellman, M. P. 2003. On market-inspired approaches to propositional satisfiability. *Artificial Intelligence* 144:125–156.
- Yates, J. 1989. *Control through Communication: The Rise of System in American Management*. Baltimore: Johns Hopkins University Press.
- Young, H. P. 1998. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton: Princeton University Press.
- Ziman, J. 2000. *Real Science: What It Is, and What It Means*. Cambridge, England: Cambridge University Press.